



## **Application Performance Management**

# **User Guide**

**Date**      2021-10-11

---

# Contents

---

<b>1 Introduction.....</b>	<b>1</b>
<b>2 Basic Concepts.....</b>	<b>2</b>
<b>3 Usage Restrictions.....</b>	<b>6</b>
<b>4 Permissions Management.....</b>	<b>9</b>
<b>5 Application Deployment.....</b>	<b>12</b>
<b>6 Dashboard.....</b>	<b>16</b>
<b>7 Alarm Center.....</b>	<b>17</b>
7.1 Viewing Alarms.....	17
7.2 Viewing Events.....	18
7.3 Setting Alarm Notification.....	19
<b>8 Topology.....</b>	<b>23</b>
<b>9 Inventory.....</b>	<b>27</b>
<b>10 Transactions.....</b>	<b>28</b>
<b>11 Tracing.....</b>	<b>30</b>
11.1 Call Chain.....	30
11.2 Method Tracing.....	31
<b>12 SQL Analysis.....</b>	<b>33</b>
<b>13 JVM Monitoring.....</b>	<b>35</b>
<b>14 Collection Management.....</b>	<b>40</b>
14.1 Agent Management.....	40
14.1.1 Installing the ICAgent.....	40
14.1.2 Upgrading the ICAgent.....	42
14.1.3 Uninstalling the ICAgent.....	42
14.2 Collection Configuration.....	43
<b>15 Configuration Center.....</b>	<b>46</b>
<b>16 FAQs.....</b>	<b>47</b>
16.1 What Data Will Be Collected and What Are They Used For?.....	47

---

16.2 How Do I Obtain an AK/SK?.....	49
16.3 How Do I Obtain an AK/SK by Creating an Agency?.....	50
<b>17 Change History.....</b>	<b>52</b>

# 1 Introduction

---

## What Is APM?

Currently, user experience has become one of core competences of applications. With the increasing application complexity and increasing number of users, application O&M faces huge challenges at normal application assurance, fast fault locating, and performance bottleneck identification.

Application Performance Management (APM) monitors and manages the performance of cloud applications in real time. APM provides performance analysis of distributed applications, helping O&M personnel quickly locate and resolve faults and performance bottlenecks.

As a cloud application diagnosis service, APM supports applications based on multiple Java frameworks. It includes powerful analytic tools, displays application status, call process, and operations performed on applications through topology views, tracing, and transactions. This helps you quickly locate faults and performance bottlenecks.

## Architecture Highlights

With the emergence of new technologies and methods, enterprises have urgent demands for fast and agile compatibility support, and need to monitor and analyze applications in multi-layer, complex, and hybrid architectures. APM provides automatic and real-time monitoring and analysis capabilities in new IT architectures such as mobile, on-cloud, and distributed systems. It supports proactive O&M and auxiliary optimization to ensure consistent user experience.

To shield the impact of technical changes on the computing and storage layers and provide stable and reliable data analysis formats, APM uses the data collection and access layer which is lightweight and irrelevant to frameworks, and supports multi-layer decoupling and independent extension. In this way, the computing, storage, and presentation layers can stably focus on analysis, calculation, and display of data.

# 2 Basic Concepts

---

## Topology

Topologies show the call and dependency relationships between applications. A topology view consists of circles, arrows, lines, and resources. Each circle represents an application, and each segment in the circle represents an instance. The fraction in each circle indicates the number of active instance/total number of instances. The data below the fraction indicates the **service latency**, calls, and errors. Each line with an arrow represents a call relationship. Thicker arrows indicate more calls. The data on arrows is throughput and **overall latency**. Throughput is the number of calls within the selected period. **Application Performance Index (Apdex)** is used to quantify user satisfaction with application performance. Different colors indicate different Apdex ranges, helping you quickly detect and locate faults.

## Transaction

A transaction is usually an HTTP request (complete process: user request > web server > database > web server > user request). In real life, a transaction is a one-time task. A user completes a task by using an application. In the example of an e-commerce application, a commodity query is a transaction, and a payment is also a transaction.

## Call Chain

By tracing and recording service calls, Application Performance Management (APM) visually restores the execution track and status of service requests in distributed systems, so that you can quickly demarcate performance bottlenecks and faults.

## Application

You can put the same type of services into an application for better performance management. For example, you can put accounts, products, and payment applications into the **Mall** application.

## Apdex

Apdex is an open standard developed by the Apdex alliance. It defines a standard method to measure application performance. The application response time is

converted into user satisfaction with application performance. The Apdex value ranges from 0 to 1.

- Apdex principles

Apdex defines the optimal threshold, T for the application response time. T is determined based on performance expectations. Based on the actual response time and T, user experience can be categorized as follows:

**Satisfied:** indicates that the actual response time is shorter than or equal to T. For example, if T is 1.5s and the actual response time is 1s, user experience is satisfied.

**Tolerating:** indicates that the actual response time is greater than T, but shorter than or equal to 4T. For example, if T is 1s, the tolerable upper threshold for the response time is 4s.

**Frustrated:** indicates that the actual response time is greater than 4T.



- Apdex calculation method

In APM, T is the threshold set in [Setting Apdex Thresholds](#), the application response latency equals to the service latency, and the Apdex value ranges from 0 to 1. The calculation formula is as follows:

Apdex = (Number of normal calls x 1 + Number of slow calls x 0.5)/Total number of calls

Among them:

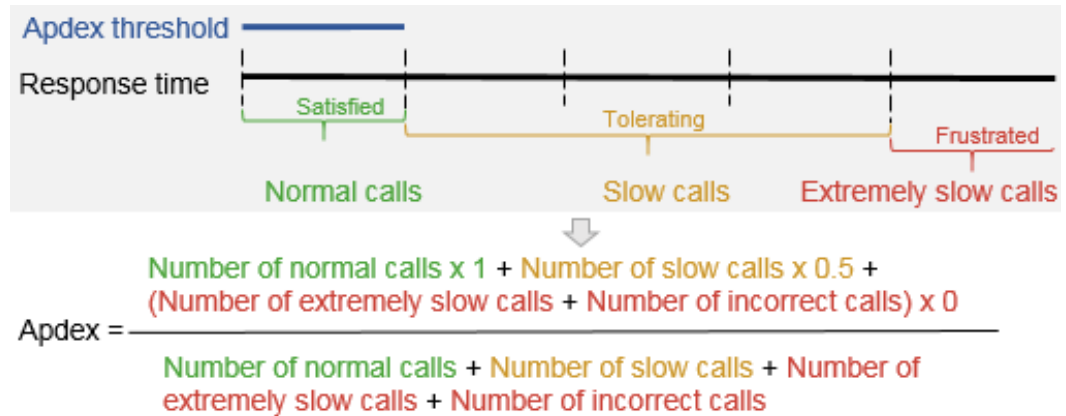
**Number of normal calls:** indicates the number of successful calls that are completed within a time period of greater than 0 but less than T.

**Number of slow calls:** indicates the number of successful calls that are completed within a time period of greater than or equal to T but less than 4T.

**Number of extremely slow calls:** indicates number of successful calls that are completed within a time period of greater than 4T.

**Total number of calls:** indicates the total number of normal calls, slow calls, extremely slow calls, and incorrect calls.

The Apdex calculation formula is as follows:



Apdex calculation results indicate application performance status, that is, user satisfaction with application performance. Apdex results are marked by different colors. For details, see [Table 2-1](#).

**Table 2-1** Apdex values

Apdex Value	Color	Description
$0.75 \leq \text{Apdex} \leq 1$	Green	Fast response; good user experience
$0.3 \leq \text{Apdex} < 0.75$	Yellow	Slow response; fair user experience
$0 \leq \text{Apdex} < 0.3$	Red	Very slow response; poor user experience

### Configuring an Apdex threshold

You can configure the Apdex threshold based on service requirements. For details, see [Setting Apdex Thresholds](#).

## TP99 Latency


TP99 latency is the minimum time meeting requirements of 99% requests. In APM, latency refers to TP99 latency.

For example, the time required for processing four requests is 10 ms, 100 ms, 500 ms, and 20 ms respectively.

In the four requests, the number of 99% requests can be calculated by multiplying 4 by 99%, and the rounding value is 4. That is, the number of 99% requests is 4. The minimum time required for the four requests is 500 ms. Therefore, TP99 latency is 500 ms.

## Overall Latency/Service Latency

Latency refers to the period from initiating a request to getting a response. In APM, the overall latency refers to the total time consumed by a request, and the service latency refers to the time consumed by a service. For example, assume that service A calls service B, and service B calls service C, as shown in the following figure:



Overall latency =  $T_A$ ; Latency of service A =  $T_A$ ; Latency of service B =  $T_{B1} + T_{B2}$ ;  
Latency of service C =  $T_C$

## Collection Probe

Probes use the bytecode enhancement technology to track calls and generate data. The data will be collected by the ICAgent and then displayed on the UI. If the memory monitoring mechanism is enabled and the instance memory usage is too high, probes enter the hibernation state and stop data collection. For details about the types of data collected by probes, see [16.1 What Data Will Be Collected and What Are They Used For?](#)

## ICAgent

ICAgent is the collection agent of APM. It runs on the server where applications are deployed to collect data obtained by the probe in real time. [Installing ICAgent](#) is a prerequisite for APM usage.



# 3 Usage Restrictions

## Supported OSs

Application Performance Management (APM) supports multiple Operating Systems (OSs). When creating an Elastic Cloud Server (ECS), select an OS supported by APM. For details, see [Table 3-1](#).

**Table 3-1** Supported OSs and versions

OS	Version	Description
SUSE	SUSE Enterprise 12 SP1 64-bit SUSE Enterprise 12 SP2 64-bit SUSE Enterprise 11 SP4 64-bit	-
openSUSE	13.2 64-bit 42.2 64-bit	-
EulerOS	2.2 64-bit	-
CentOS	7.4 64-bit 7.3 64-bit 7.2 64-bit 7.1 64-bit 6.9 64-bit 6.8 64-bit 6.5 64-bit 6.3 64-bit	-

OS	Version	Description
Ubuntu	14.04 server 64-bit 16.04 server 64-bit	-
CoreOS	10.10.5 64-bit	-
Fedora	24 64-bit	The 25 64-bit version has been planned and is being tested.
Debian	To be supported	The 7.5.0 32-bit and 7.5.0 64-bit versions have been planned and are being tested.

## Supported Types

Currently, APM can connect to only Java applications. APM supports mainstream Java frameworks, web servers, communications protocols, and databases. For details about the supported types, see [Table 3-2](#).

**Table 3-2** Supported types

Type	Name	Version
Tool	JDK	JDK 7 and JDK 8
Communications protocol	HTTP client	Apache HttpClient 3, Apache HttpClient 4, and JDK HttpURLConnection
Java framework	CXF Client	2.6.0–3.2.1
	iBatis	2.3.0 and 2.3.4.726
	Jersey	2.0–2.9.1
	MyBatis	1.0.0–1.3.1 (MyBatis-Spring) and 3.0.1–3.4.5 (MyBatis 3)
	Spring	3.1.x–5.0.x
	Spring Boot	1.2.x–1.5.x
	Dubbo	2.5.3–2.5.4 (Dubbo RPC and Dubbo REST)
	CSE	0.4–0.5 (REST over Servlet, REST over Vertx, and Highway RPC)
Database	MySQL	5.1.x
	Oracle	ojdbc5, ojdbc6, and ojdbc14
	Sybase	2.6.0–3.2.1

Type	Name	Version
	MariaDB	1.3.x
	VoltDB	6.x-7.x
	PostgreSQL	9.0.x, 9.1.x, 9.2.x, 9.3.x, 9.4.x, 42.0.x, and 42.1.x
Web server	Tomcat	6.x, 7.x, and 8.x
	Jetty	7.6.x-8.0.0 and 8.1.x-9.x.x
	JBoss	7.0.0-7.1.3 and 7.2.0
	Undertow	1.4.x
Message queue	ActiveMQ	5.6.x-5.15.x
	RocketMQ	4.1.x-4.2.x
	RabbitMQ	1.3.3 and later (spring-rabbit), 2.7.x (amqp-client), 2.6.0, and 3.6.5
	Kafka	0.9.0.1-0.10.0.2
NoSQL	Redis	2.0.0-2.9.0
	Memcache	2.9.0-2.12.3 (Arcus)
	MongoDB	3.0.x-3.6.x
	Cassandra	2.1.x-3.2.x
	ZooKeeper	1.0.x (com.github.adyliu.zkclient) and 0.1.x (com.github.sgroscupf.zkclient)
	ElasticSearch	2.4.x and 5.1.x
REST Client	Common HTTP	2.x, 3.x, 4.x (HttpClient), and ALL (URLConnection)

 **NOTE**

More types are being developed.

# 4 Permissions Management

---

If you need to assign different permissions to employees in your enterprise to access your Application Performance Management (APM) resources, Identity and Access Management (IAM) is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your cloud resources.

With IAM, you can use your account to create IAM users for your employees, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use APM resources but must not delete them or perform any high-risk operations. To achieve this result, you can create IAM users for the software developers and grant them only the permissions required for using APM resources.

If your account does not need individual IAM users for permissions management, you may skip over this chapter.

IAM can be used free of charge. You pay only for the resources in your account.

## APM Permissions

By default, new IAM users do not have any permissions assigned. You need to add a user to one or more groups, and assign permissions policies or roles to these groups. The user then inherits permissions from the groups it is a member of. This process is called authorization. After authorization, the user can perform specified operations on APM.

APM is a project-level service deployed and accessed in specific physical regions. To assign APM permissions to a user group, specify the scope as region-specific projects and select projects for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing APM, the users need to switch to a region where they have been authorized to use this service.

[Table 4-1](#) lists all the system permissions supported by APM.

**Table 4-1** System permissions supported by APM

Role	Description	Category
APM FullAccess	Full permissions for APM	System-defined policy
APM ReadOnlyAccess	Read-only permissions for APM	System-defined policy
APM Administrator	Full permissions for APM	System-defined role

**Table 4-2** lists the common operations supported by each system-defined policy or role of APM. Choose appropriate policies or roles as required.

**Table 4-2** Common operations supported by each system-defined policy or role of APM

Operation	APM FullAccess	APM ReadOnlyAccess	APM Administrator
Obtaining application topology information	√	√	√
Modifying application topology configuration	√	x	√
Deleting application topology configuration	√	x	√
Adding application topology configuration	√	x	√
Obtaining slow SQL analysis results	√	√	√
Obtaining tracing data	√	√	√
Updating tracing configuration	√	x	√
Querying APM configuration	√	√	√

<b>Operation</b>	<b>APM FullAccess</b>	<b>APM ReadOnlyAccess</b>	<b>APM Administrator</b>
Adding APM configuration	√	x	√
Deleting APM configuration	√	x	√
Querying the ICAgent list	√	√	√
Installing the ICAgent	√	x	√
Querying the ICAgent version	√	√	√
Upgrading the ICAgent version	√	x	√
Uninstalling the ICAgent	√	x	√
Delivering an ICAgent event	√	x	√

# 5 Application Deployment

---

This section describes how to deploy applications for performance management.

You need to perform operations based on application deployment modes. Currently, Application Performance Management (APM) supports application deployment through:

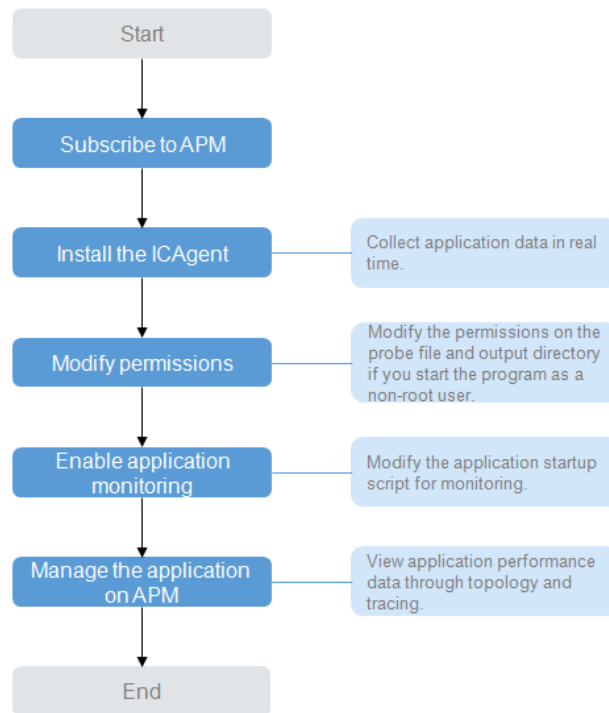
- Elastic Cloud Server (ECS). For details, see [ECS Mode](#).
- ServiceStage. For details, see [ServiceStage Mode](#).
- Application Orchestration Service (AOS). For details, see [AOS Mode](#).
- Cloud Container Engine (CCE). For details, see [CCE Mode](#).

## ECS Mode

### Prerequisites

1. You have created an ECS server.
2. The ECS server meets the requirements in [Supported OSs](#).
3. The ECS server meets the requirements in [Supported Types](#).
4. The time and time zone of the local browser are consistent with those of the ECS server.

### Process



**Step 1: Install the ICAgent on a VM**

**(Optional) Step 2: Modify Permissions**

If you start the program as a non-root user, log in to the ECS server and run the following commands to modify the permissions on the probe file and output directory before enabling application monitoring:

```

chmod -R 777 /opt/oss/servicemgr/ICAgent/pinpoint/
mkdir -p /paas-apm/collectors/pinpoint
chmod -R 777 /paas-apm
    
```

**Step 3: Configure the Application Startup Script and Restart the Application**

1. On the ECS server, add configuration items in the following table after the **java** keyword in the Java application startup script to ensure that the Java application is monitored by APM.

**Table 5-1** Configuration items to be added

Parameter	Description
-javaagent	JAR package that collection probes depend on. The fixed value is <b>/opt/oss/servicemgr/ICAgent/pinpoint/pinpoint-bootstrap.jar</b> .
-Dapm_application	Application name. The value must be 1 to 64 characters starting with a letter or an underscore (_). Only lowercase letters, digits, hyphens (-), and underscores are allowed.



Parameter	Description
-Dapm_tier	Name of the application microservice. The value must be 1 to 64 characters starting with a letter or an underscore (_). Only lowercase letters, digits, hyphens (-), and underscores are allowed.

- Execute the modified application startup script to enable application monitoring.

 **NOTE**

**Example of the modified startup script**

The following shows an example startup script of the **Vmall** application with the **vmall-dao-service** and **vmall-user-service** services. You need to configure your script as required.

- Original startup script:  

```
java -Xmx512m -jar /root/testdemo/ecommerce-persistence-service-0.0.1-SNAPSHOT.jar --spring.config.location=file:/root/testdemo/application_dao.yml > dao.log 2>&1 &
java -Xmx512m -jar /root/testdemo/ecommerce-user-service-0.0.1-SNAPSHOT.jar --spring.config.location=file:/root/testdemo/application_userservice.yml > user.log 2>&1 &
```
- Modified startup script (differences are in bold):  

```
java -javaagent:/opt/oss/servicemgr/ICAgent/pinpoint/pinpoint-bootstrap.jar -Dapm_application=vmall -Dapm_tier=vmall-dao-service -Xmx512m -jar /root/testdemo/ecommerce-persistence-service-0.0.1-SNAPSHOT.jar --spring.config.location=file:/root/testdemo/application_dao.yml > dao.log 2>&1 &
java -javaagent:/opt/oss/servicemgr/ICAgent/pinpoint/pinpoint-bootstrap.jar -Dapm_application=vmall -Dapm_tier=vmall-user-service -Xmx512m -jar /root/testdemo/ecommerce-user-service-0.0.1-SNAPSHOT.jar --spring.config.location=file:/root/testdemo/application_userservice.yml > user.log 2>&1 &
```

**Step 4: Manage the Application on APM**

After the applications run for about three minutes, its data will be displayed on the APM console. You can log in to the APM console and optimize application performance through topology and tracing. For details, see later chapters.

**ServiceStage Mode**

ServiceStage is the one-stop DevOps platform service oriented for enterprises and developers. If you select probes when using ServiceStage to create or release applications, APM is automatically connected to the applications. After the applications run for about three minutes, you can view the application information on the **Topology** and **Transactions** pages of APM.

**AOS Mode**

For AOS, when you add the designer Pinpoint to templates during compilation, APM collection probes are added to stacks. After templates are compiled and stacks are created, APM is automatically connected to stack applications. After the stacks run for about three minutes, you can view the application information on the **Topology** and **Transactions** pages of APM.

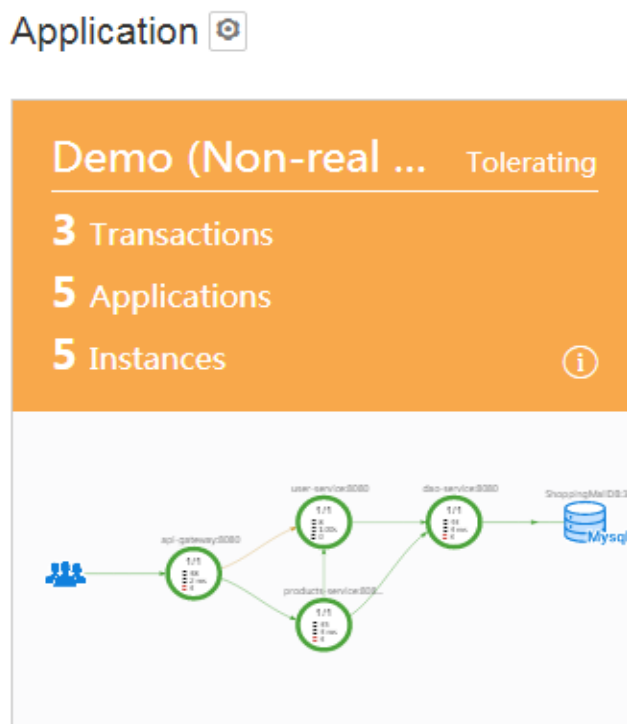
## CCE Mode

CCE provides container application management services. If you select probes when creating or upgrading applications, APM collection probes are installed on the applications. After the applications run for about three minutes, you can view the application information on the **Topology** and **Transactions** pages of APM.

# 6 Dashboard

You can quickly learn about the health status of applications through the dashboard.

**Figure 6-1** Dashboard page



You can delete a service card in the following scenarios:

- The service connected to Application Performance Management (APM) has been deleted.
- The ICAgent has been uninstalled and service data does not need to be collected.

If the service connected to APM is still running, the service card will be displayed again three minutes later upon deletion.

# 7 Alarm Center

---

- [7.1 Viewing Alarms](#)
- [7.2 Viewing Events](#)
- [7.3 Setting Alarm Notification](#)


## 7.1 Viewing Alarms

Alarms are reported when Application Performance Management (APM) or an external service, such as Application Orchestration Service (AOS), ServiceStage, or Cloud Container Engine (CCE), is abnormal or may cause exceptions. Alarms need to be handled. Otherwise, service exceptions may occur.

### Procedure

- Step 1** Log in to the APM console.
- Step 2** In the navigation pane, choose **Alarm Center > Alarm List**.
- Step 3** View alarms on the **Alarm List** page.
  1. Set a time range to view alarms. There are two methods to set a time range:  
Method 1: Use the predefined time label, for example, **Last 1 hour**, **Last 6 hours**, or **Last 1 day**. You can choose one based on service requirements.  
Method 2: Specify the start time and end time to customize a time range. The time range can be 30 days at most.
  2. Set filter criteria and click **Search** to view alarms.  
Click **Reset** to reset filter criteria.
- Step 4** You can also perform the operations described in [Table 7-1](#).

**Table 7-1** Operations

Operation	Method	Description
Viewing alarm statistics	View alarm statistics that meet filter criteria within a specific time range through a bar graph.	-
Clearing alarms	In the current alarm list, click <b>Clear</b> in the <b>Operation</b> column for a target alarm.	<ul style="list-style-type: none"> <li>You can click <b>Clear</b> in the <b>Operation</b> column after the problem is resolved.</li> <li>Alarms that are cleared cannot be queried.</li> </ul>
Viewing alarm details	Click <b>View Details</b> in the <b>Operation</b> column to view alarm details.	-
Viewing the latest alarms	Click  on the right of the page to view the latest three alarms.	-

----End

## 7.2 Viewing Events

Events generally carry some important information. They are reported when Application Performance Management (APM) or an external service, such as Application Orchestration Service (AOS), ServiceStage, or Cloud Container Engine (CCE) encounters some changes. Such changes do not necessarily cause service exceptions. Events do not need to be handled.

### Procedure

**Step 1** Log in to the APM console.

**Step 2** In the navigation pane, choose **Alarm Center > Event List**.

**Step 3** View events on the **Event List** page.

- Set a time range to view events. There are two methods to set a time range:
  - Method 1: Use the predefined time label, for example, **Last 1 hour**, **Last 6 hours**, or **Last 1 day**. You can choose one based on service requirements.
  - Method 2: Specify the start time and end time to customize a time range. The time range can be 30 days at most.
- Set filter criteria and click **Search** to view events.  
Click **Reset** to reset filter criteria.

**Step 4** You can also perform the operations described in [Table 7-2](#).

**Table 7-2** Operations

Operation	Method	Description
Viewing event statistics	View event statistics that meet filter criteria within a specific time range through a bar graph.	-

----End

## 7.3 Setting Alarm Notification

Application Performance Management (APM) supports alarm notification. That is, a certain type of alarms can be sent to specified users by emails. In this way, they can identify and rectify cluster exceptions at the earliest time, preventing service loss.

You can create a maximum of 10 notification rules. If the number of notification rules reaches 10, delete unnecessary notification rules.

If no notification rules exist, you will not receive any alarm notifications. In that case, you can only view alarms by logging in to the APM console and choosing **Alarm Center > Alarm List** in the navigation pane.

Currently, APM supports creation of notification rules only for the alarms listed in [Table 7-3](#).

**Table 7-3** Alarm types

Alarm Type	Description
Collector installation alarm	This alarm is generated when the ICAgent fails to be installed, upgraded, or uninstalled, or is abnormal.

 **NOTE**

More types of alarms are being developed.

### Creating a Notification Rule

- Step 1** Log in to the APM console.
- Step 2** In the navigation pane, choose **Alarm Center > Notification Rules**, and click **Create Notification Rule**.
- Step 3** Create a topic, configure a topic policy, and add subscribers to the topic. If you have configured them, skip the following operations.
  1. Access the SMN console: When APM is interconnected with Simple Message Notification (SMN), click **Create SMN Topic** to access the SMN console.

2. Create a topic: In the navigation pane of the SMN console, choose **Topic Management > Topics**. Then click **Create Topic**. On the page that is displayed, enter a topic name and click **OK**.
3. Configure a topic policy and add subscribers: On the **Topics** page of the SMN console, choose **More > Configure Topic Policy** in the **Operation** column of the target topic. On the page that is displayed, configure a topic policy according to **Figure 7-1**. Otherwise, notifications may fail to be sent. Then, add subscribers, that is, email receivers of notifications. In this way, when an exception occurs in a cluster, APM can broadcast the alarm information to the subscribers in real time.

**Figure 7-1** Configuring a topic policy

**Configure Topic Policy** ×

Topic Name test321

Policy **Basic** | Advanced

Users who can publish messages to this topic

Topic creator

All users

Specified user accounts

Enter one or more domain IDs or URNs. Each line must contain only one domain ID or URN.

For details about how to obtain the account ID, click [Learn how](#)

Services that can publish messages to this topic

OBS

apm

**OK** Cancel

**Step 4** Enter a rule name, select an alarm type (for details, see **Table 7-3**), select a created topic in **Step 3**, select a cluster to be monitored, and click **Create**, as shown in **Figure 7-2**.

After the notification rule is created, if an alarm that meets the notification rule is generated, APM automatically sends notifications by emails.

**Figure 7-2** Creating a notification rule

## Create Notification Rule

Rule Name

Alarm Type  ▼

Alarms that are similar are classified under the same type.

Topic  x ▼

[Create SMN Topic](#)

Ensure that APM is selected as one of the services that can publish messages to the topic when configuring a topic policy on the SMN page. For details, see [Configuring Topic Policies](#).

Cluster  ▼ [Create Cluster](#)

**NOTE**

If the message **Sorry, you do not have the permission to access Simple Message Notification (SMN)**. is displayed when you select a topic, it is because you have logged in to APM as an Identity and Access Management (IAM) user who does not have the permission to access SMN. Contact the administrator (account to which the IAM user belongs) to add the SMN access permission.

----End


## More Operations

After creating a notification rule, you can also perform the operations described in [Table 7-4](#).

**Table 7-4** Related operations

Operation	Description
Modifying a notification rule	Click <b>Modify</b> in the <b>Operation</b> column.
Deleting a notification rule	<ul style="list-style-type: none"> <li>To delete one notification rule, click <b>Delete</b> in the <b>Operation</b> column.</li> <li>To delete one or more notification rules, select it or them and click <b>Delete</b> above the notification rule list.</li> </ul>



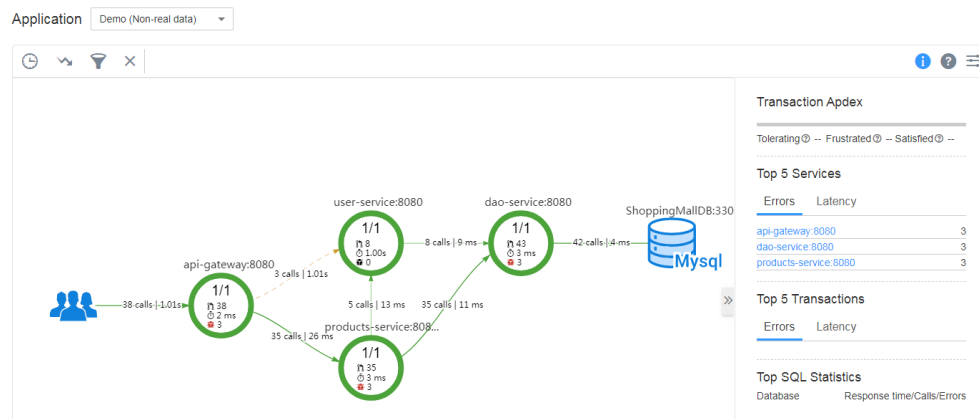
Operation	Description
Searching for a notification rule	Enter a keyword of the notification rule name in the search box in the upper right corner and click  .

# 8 Topology

In a topology, each circle represents a service, each segment of a circle represents an instance, and each arrow represents a call relationship. In addition, Application Performance Management (APM) can display the call relationships between applications. Each circle can also represent an application. When a circle represents an application, right-click a circle and choose **View Application** to go to the associated application topology page.

Different colors on the circle represent different health statuses of instances. The color is determined by the **Application Performance Index (Apdex)** value. If an Apdex value is closer to **1**, the corresponding application is healthier.

## Topology Page




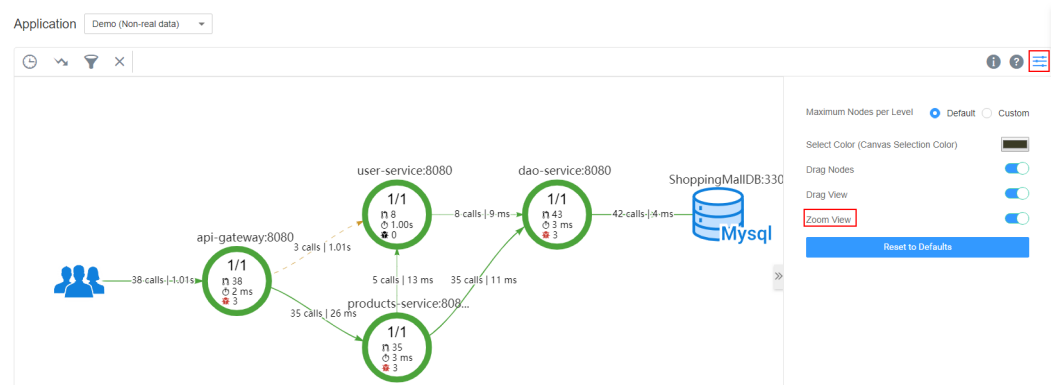
1. **Table 8-1** provides topology description.

**Table 8-1** Topology description

Color	Instance	Call
Green	$0.75 \leq \text{Apdex} \leq 1$ The instance responds quickly when it is called.	$0.75 \leq \text{Apdex} \leq 1$ Quick response.

Color	Instance	Call
Yellow	0.3 ≤ Apdex < 0.75 The instance responds slowly when it is called.	0.3 ≤ Apdex < 0.75 Slow response.
Red	0 ≤ Apdex < 0.3 The instance responds very slowly when it is called.	0 ≤ Apdex < 0.3 Very slow response.
Gray	The instance is not called.	N/A
Black	The instance has been deleted.	N/A

- On the **Topology** page, click  to configure the topology. For example, if **Zoom View** is disabled, you cannot zoom in or out the topology.

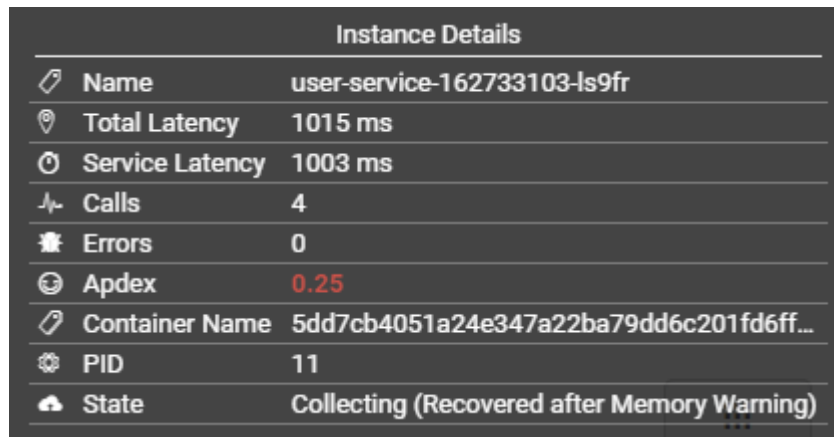


## Locating Problems Based on the Topology

The following describes how to locate an instance with a slow response:

- Step 1** Log in to the APM console.
- Step 2** In the navigation pane, choose **Topology**.
- Step 3** On the **Topology** page, select the time range during which a problem occurred in the upper right corner.
- Step 4** In the topology, view the instance (highlighted in red) with a slow response, as shown in [Figure 8-1](#).

**Figure 8-1** Abnormal instance



Instance Details	
Name	user-service-162733103-ls9fr
Total Latency	1015 ms
Service Latency	1003 ms
Calls	4
Errors	0
Apdex	0.25
Container Name	5dd7cb4051a24e347a22ba79dd6c201fd6ff...
PID	11
State	Collecting (Recovered after Memory Warning)

- Step 5** (Optional) For the service containing multiple instances, right-click an instance and choose **Expand** from the shortcut menu to view the call relationships between the instances to preliminarily identify the abnormal instance.
- Step 6** Right-click the instance and choose **Find Call-Chain** from the shortcut menu. On the **Call Chain** page that is displayed, further locate the problem based on call duration and other parameters.

----End

## Configuring Transaction Apdex Thresholds

The response time of different transactions is different. APM enables you to configure different Apdex thresholds for transactions. For example, if a login takes more than 50 ms, the response is slow. If a transaction query takes more than 10 ms, the response is slow. In this case, you need to set different Apdex thresholds for the login and query transactions.

- Step 1** Log in to the APM console.
- Step 2** In the navigation pane, choose **Topology**.
- Step 3** On the topology page, move the mouse cursor over the circle diagram, right-click it, and click **Edit Threshold**.
- Step 4** Modify the transaction Apdex threshold and click **Apply**.

×

### Edit Tier - Apdex Threshold (ms)

cusotmer-service:9001

					Total Call:	Apdex	Apdex T	Current Apdex
	<span style="color: purple;">●</span>							
	GET_/hello/undertow/{name}							
	6	0	0	0	6	1	100	100
	ALL							
	6	0	0	0	6	1	100	100
	<b>Note: Change Will Be Apply For Only New Snapshots</b>							

ApplyCancel

----End

# 9 Inventory

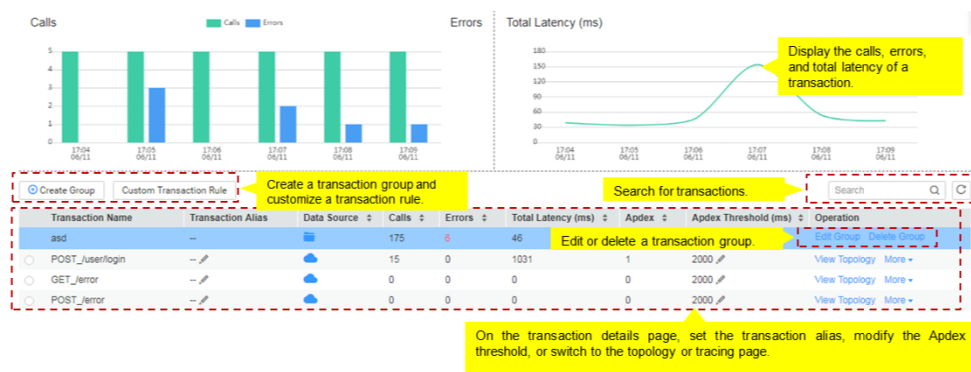
---

On the **Inventory** page, application details (such as application service types and resource IDs) are displayed, helping you locate problems.

# 10 Transactions

A transaction may require multiple calls between services. Any slow or incorrect call may lead to slow responses. During routine O&M, you can analyze the transactions with slow responses to locate and solve application problems, thereby improving user experience of services.

## Transaction Page



## Analyzing Problems Based on Transactions

The following describes how to locate the cause of a transaction with an extremely slow response.

- Step 1** Log in to the Application Performance Management (APM) console.
- Step 2** In the navigation pane, choose **Transactions**.
- Step 3** In the transaction list, select a transaction with an extremely slow response.

The screenshot shows a table with columns: Transaction Name, Transaction Alias, Data Source, Calls, Errors, Total Latency (ms), Apex, Apex Threshold (ms), and Operation. The transaction 'ALL\_/user/\*\*' is highlighted in blue, indicating it is selected. It has 75 calls, 0 errors, and a total latency of 1023 ms. The Apex value is 1, and the Apex Threshold is 2000 ms.

Transaction Name	Transaction Alias	Data Source	Calls	Errors	Total Latency (ms)	Apex	Apex Threshold (ms)	Operation
ALL_/product/**	--		875	33	42	0.96	2000	View Topology More
ALL_/user/**	--		75	0	1023	1	2000	View Topology More

- Step 4** Click **View Topology** in the **Operation** column to view the topology of the transaction and view the instance details in the topology.

**Step 5** In the transaction topology, right-click the instance with an extremely slow response and choose **Find Call-Chain** from the shortcut menu. On the **Call Chain** page that is displayed, further locate the problem based on call duration and other parameters.

----End



# 11 Tracing

---

[11.1 Call Chain](#)

[11.2 Method Tracing](#)

## 11.1 Call Chain

By tracing and recording service calls, Application Performance Management (APM) visually restores the execution track and status of service requests in distributed systems, so that you can quickly demarcate performance bottlenecks and faults.

### Demarcating a Performance Bottleneck

- Step 1** Log in to the APM console.
  - Step 2** In the navigation pane, choose **Tracing > Call Chain**.
  - Step 3** At the top of the **Call Chain** page, select the time range, application name, and service name from three drop-down lists, and click **Search** to query the corresponding call chains.
  - Step 4** (Optional) On the **Call Chain** page, click **Advanced** in the upper right corner, specify filter criteria, and click **Search** to search for the desired call chain.
  - Step 5** In the **Operation** column, click **View Call Relationship**.
  - Step 6** Check the value in the **Time Line (ms)** column to locate the method with long duration and identify the performance bottleneck.
  - Step 7** (Optional) View additional information to further locate the cause.  
Click **View Details** in the **Operation** column to view detailed call information.
- End

### Locating Problems Based on Additional Information

- Step 1** Log in to the APM console.

- Step 2** In the navigation pane, choose **Tracing > Call Chain**.
- Step 3** At the top of the **Call Chain** page, select the time range, application name, and service name from three drop-down lists, and click **Search** to query the corresponding call chains.
- Step 4** (Optional) On the **Call Chain** page, click **Advanced** in the upper right corner, specify filter criteria, and click **Search** to search for the desired call chain.
- Step 5** In the **Status** column, check the application status and find out the faulty service.
- Step 6** Click **View Call Relationship**, check whether the return value is normal, and locate the fault.
- Step 7** (Optional) View additional information to further locate the cause.  
Click **View Details** in the **Operation** column to view detailed call information.  
----End

## 11.2 Method Tracing

Method tracing is used to dynamically trace a method of a class. When the method of this class is called, Application Performance Management (APM) collects the call data of the method based on configured method tracing rules by using probes, and displays the call data on the **Call Chain** page. Method tracing is designed for application developers to locate method-level performance problems online.

APM traces APIs released by third-party open-source components, but may not trace specific methods of your applications. To monitor important methods in applications or methods of third-party open-source components that are not supported by APM, you need to customize method tracing. After the configuration is complete, you can view the call data of the method on the **Call Chain** page.

- Step 1** Log in to the APM console.
- Step 2** In the navigation pane, choose **Tracing > Method Tracing**.
- Step 3** Customize a method tracing rule and start method tracing.

Specifically, on the **Method Tracing** page, click **Add Method Tracing Rule**, set the parameters according to the following figure, and click **OK**.

### Add Method Tracing Rule

Service Name	<input type="text" value="user-service"/>
* Class Name	<input type="text" value="com.██████████.api.UserController"/>
* Method Name	<input type="text" value="Login"/>
Method Parameter	<input type="text" value="com.██████████.entity.User user,java.lang.Integer workload"/>
Value	<input type="text"/>
Collect Method Stack Info <input type="checkbox"/>	Collect All Matched Call Info <input type="checkbox"/>

 **NOTE**

- If **Method Parameter** is not set, the methods of the same method name are used for collection by default.
- If **Value** is not set, the values of the method are not filtered during collection.
- If **Collect Method Stack Info** is enabled, the method stack information is collected.
- If **Collect All Matched Call Info** is enabled, all matched tracing method information is collected. If this function is disabled, tracing method information is collected based on the sampling ratio (common or intelligent sampling) set in [14.2 Collection Configuration](#).

**Step 4** Preliminarily locate service performance problems based on the total call duration and call status displayed at the bottom of the page.

**Step 5** Click **View Call Relationship** in the **Operation** column to view the method-level call relationships.

----End

# 12 SQL Analysis

---

Application Performance Management (APM) displays key metrics, including database, SQL statement calls, latency, and error calls for analyzing database performance problems caused by abnormal SQL statements. Abnormal SQL statements include slow and error SQL statements.

APM supports multiple mainstream databases, including Cassandra, Memcached, MongoDB, MySQL, Oracle, PostgreSQL, and Redis. Databases here include the relational databases that use SQL, including MySQL, Oracle, and PostgreSQL, but do not include non-relational databases, such as Cassandra, Memcached, MongoDB, and Redis.

## Analyzing Abnormal SQL Statements

When an SQL statement of a database is abnormal, performance problems such as service timeout may occur. During routine O&M, you can monitor key metrics, such as the error duration and latency of databases, locate the SQL statements that take a long time to execute, operate at low efficiency, or fail to be called, and then analyze and optimize them.

The SQL analysis function determines whether to collect SQL data. Before performing the following steps, ensure that this function is enabled. Otherwise, no SQL data can be queried. This function is enabled by default. If it is disabled, choose **Agent > Configuration** in the navigation pane and then enable it.

- Step 1** Log in to the APM console.
- Step 2** In the navigation pane, choose **SQL Analysis**.
- Step 3** On the **SQL Analysis** page, select the time range during which a problem occurred in the upper right corner.
- Step 4** On the **Overview** tab page, locate the faulty database in the application based on key database metrics. If a database has long response time and many call errors, performance problems may occur.
- Step 5** Analyze the performance problem of the database.

Click the **SQL Analysis** tab, and locate the abnormal SQL statement in the SQL statement list.

**Step 6** Analyze the cause of the abnormal SQL statement.

1. Click the abnormal SQL statement to go to the **Call Chain** page and check the impact of the abnormal SQL statement on the entire service.
2. Click **View Call Relationship** in the **Operation** column to find the method of the abnormal SQL statement. Analyze the cause of the abnormal SQL statement in this method. For example, no index is not used, the data volume is overlarge, the syntax is incorrect, or a deadlock occurs. After finding out the cause, optimize the SQL statement accordingly.

----End

---

# 13 JVM Monitoring

---

JVM monitoring displays the memory and thread of the operating environment for Java applications. You can monitor metric trends in real time to analyze performance.

On the **Memory** and **Thread** tab pages, you can view the memory and thread graphs to quickly locate problems such as memory leakage and thread exceptions.

## Memory Graphs

As shown in [Figure 13-1](#), in a selected time range, the trends of the maximum, committed, and used memory in different JVM memory spaces (such as the total memory, heap memory, and non-heap memory spaces) of an instance are displayed. In addition, the Garbage Collection (GC) duration and times are also displayed.

**Figure 13-1** Memory graphs



### JVM memory

JVM memory consists of heap and non-heap memory.

- **Heap memory:** A heap is the data area where the JVM is running. It allocates memory for all class instances and arrays. Heap memory of objects is reclaimed by the automatic memory management system called garbage collector. Heap space consists of eden space, survivor space, and tenured space.
- **Non-heap memory:** Memory (excluding heap memory) managed by JVM. Non-heap space consists of code cache and permanent space (or meta space).

Java heap is the main space of garbage collector management, also called **Garbage Collection Heap**. GC mode includes **Full GC** and **Minor GC**.

**Table 13-1** Memory spaces

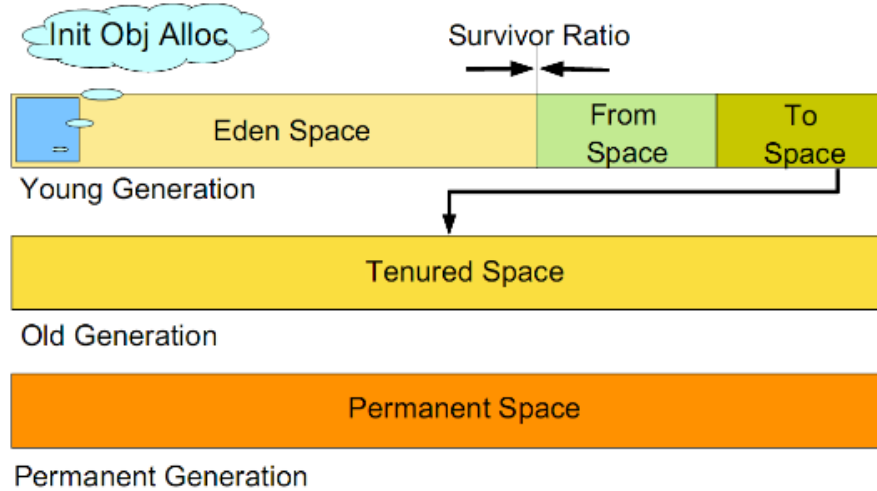
Space Name	Description
Eden space	Initially allocates memory from the thread pool to most objects.
Survivor space	Stores the eden space's objects that are not reclaimed during GC.
Tenured space	Maintains objects that have been stored in the survivor space for a period of time.
Code cache	Compiles and stores local code.
Permanent space	Stores static data of VMs, for example, class and method objects.
Meta space	Stores local class metadata. In versions later than Java 8, permanent space is replaced by meta space.
Full GC	Indicates the GC performed in the entire heap space (covering young-, old-, and permanent-generation spaces) when the memory space does not meet allocation requirements after memory reclamation.
Minor GC	Indicates the GC performed in the young-generation space (including eden and survivor spaces) when the allocated memory is insufficient.

JVM collects garbage based on generations. JVM heap space is divided into old- and young-generation spaces. More than 90% objects that exist only for a short period of time are stored in the young-generation space, while objects that have long life cycles are stored in the old-generation space. Young-generation space is further divided into eden space and two survivor spaces. New objects are initially allocated to the eden space. The survivor spaces are used as the buffer between eden space and tenured space. Objects that are survived after several rounds of



GC in the survivor spaces are then transferred to the old-generation space, as shown in [Figure 13-2](#).

**Figure 13-2** Memory spaces



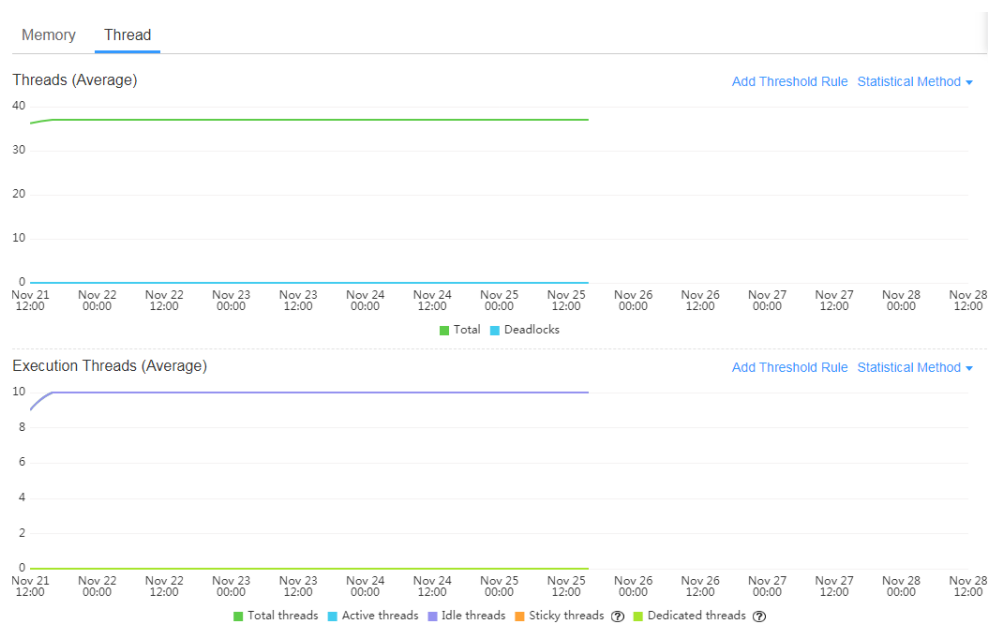
**NOTE**

There are two survivor spaces, which are represented by **from** and **to** pointers. The **to** pointer points to the empty survivor space.

## Thread Graphs

As shown in [Figure 13-3](#), in a selected time range, the trends of total threads, sticky threads, dedicated threads, and other threads are displayed.

**Figure 13-3** Thread graphs



**Table 13-2** Thread description

Thread Name	Description
Total threads	Both active and standby threads are included. Sticky threads and dedicated threads become standby threads after being executed.
Deadlock threads	When two or more threads encounter resource conflicts or abnormal communication, the system enters the deadlock state.
Sticky threads	If the time taken to process a request by a thread exceeds the preset maximum time, the thread is called a sticky thread.
Dedicated threads	If the time taken to process a request by a thread exceeds the normal execution time but does not exceed the maximum time of a sticky thread, the thread is called a dedicated thread.
Total executed threads	Both active and idle threads are included.
Active threads	Sticky threads, dedicated threads, and threads that are being executed are included.
Idle threads	Threads are in idle state. When there is no task, a thread is in the idle state. When receiving a request, the thread pool assigns an idle thread to the request. After the assigned task is completed, the idle thread returns to the thread pool and waits for another task.

# 14 Collection Management

---

[14.1 Agent Management](#)

[14.2 Collection Configuration](#)

## 14.1 Agent Management

### 14.1.1 Installing the ICAgent

#### Installing the ICAgent on a VM

Before installing the ICAgent, ensure that the time and time zone of the local browser are consistent with those of the server. If multiple servers are deployed, ensure that the local browser and multiple servers use the same time zone and time. Otherwise, metric data of applications and servers displayed on the console may be incorrect.

**Step 1** [Obtain the Access Key ID/Secret Access Key \(AK/SK\)](#).

**Step 2** Log in to the Application Performance Management (APM) console. In the navigation pane, choose **Agent > Management**. Choose the desired host from the drop-down list on the right.

**Step 3** Click **Install ICAgent**.

**Step 4** Generate the ICAgent installation command and copy it.

1. Enter the obtained AK/SK in the text box to generate the ICAgent installation command.

 **NOTE**

Ensure that the AK/SK are correct. Otherwise, the ICAgent cannot be installed.

2. Click **Copy Command**.

## Install ICAgent

Installation Mode

Obtain AK/SK

Create Agency

You can install ICAgent in either of the above ways. If you have installation for multiple hosts, please refer to [Inherited Batch Installation](#).

**1** Enter the AK/SK to generate the installation command. [How to Obtain an AK/SK?](#)

AK

SK

**2** Copy Command

Command Generated [Copy Command](#)

```
curl http://icagent-cn-cmcc1.obs-cn-cmcc1.cmyun.cn/ICAgent_linux/apm_agent_install.sh > apm_agent_install.sh && REGION=cn-cmcc1 bash apm_agent_install.sh
```

**Step 5** Use a remote login tool to log in to the server where the ICAgent is to be installed as the **root** user and run the command copied in [Step 4.2](#) to install the ICAgent.

### NOTE

- If the message **ICAgent install success** is displayed, the ICAgent is successfully installed in the `/opt/oss/servicemgr/` directory. After the ICAgent is successfully installed, choose **Agent > Management** in the navigation pane to view the ICAgent status.
- If the installation fails, uninstall the ICAgent and then install it again. If the problem persists, contact technical support.

----End

## Installing the ICAgent on a CCE Cluster

**Step 1** In the navigation pane, choose **Agent > Management**.

**Step 2** Choose the desired cluster from the drop-down list on the right.

**Step 3** Click **Install ICAgent** in the upper left corner of the page. In the **Install ICAgent** dialog box, click **Yes**.

Agent Management

Node Name	Node IP Address	ICAgent Status	ICAgent Version	Updated At
		Uninstall	--	--
		Uninstall	--	--

The ICAgent will be installed on all hosts in the cluster.

----End

## 14.1.2 Upgrading the ICAgent

To ensure better collection experience, Application Performance Management (APM) will continuously upgrade ICAgent versions. When the system displays a message indicating that a new ICAgent version is available, perform the following operations:

 **NOTE**

If the ICAgent has a critical bug, the system will upgrade the ICAgent version.

**Step 1** Log in to the APM console.

**Step 2** In the navigation pane, choose **Agent > Management**.

**Step 3** Select **Cluster: XXX** or **Other: user-defined nodes** from the drop-down list on the right of the page.

**Step 4** Upgrade the ICAgent.

- If you select **Cluster: xxx** in **Step 3**, directly click **Upgrade ICAgent** to upgrade the ICAgent on all hosts in the cluster at a time.
- If you select **Other: user-defined nodes** in **Step 3**, select a desired host and then click **Upgrade ICAgent**.

**Step 5** In the displayed **Upgrade ICAgent** dialog box, click **Yes**. The ICAgent begins to be upgraded. This operation takes about 1 minute to complete. When the ICAgent status changes from **Upgrading** to **Running**, the ICAgent is successfully upgraded.

----End

## 14.1.3 Uninstalling the ICAgent

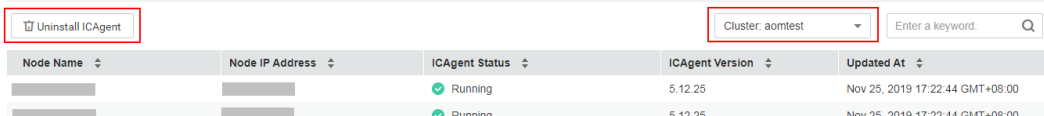
If the ICAgent on the server is uninstalled, server O&M will be affected, making topology and tracing functions unavailable. Exercise caution when performing this operation.

**Step 1** Log in to the Application Performance Management (APM) console. In the navigation pane, choose **Agent > Management**.

**Step 2** Choose the desired cluster from the drop-down list on the right.

**Step 3** Click **Uninstall ICAgent** in the upper left corner of the page. In the **Uninstall ICAgent** dialog box, click **Yes**.

Agent Management ©



Node Name	Node IP Address	ICAgent Status	ICAgent Version	Updated At
		Running	5.12.25	Nov 25, 2019 17:22:44 GMT+08:00
		Running	5.12.25	Nov 25, 2019 17:22:44 GMT+08:00

The ICAgent will be uninstalled from all the hosts in the cluster.

----End

## 14.2 Collection Configuration

To reduce memory, database, and disk space usage, you can implement collection configuration as required. The collection settings take effect for all applications you selected.

### Procedure


**Step 1** Log in to the Application Performance Management (APM) console.

**Step 2** In the navigation pane, choose **Agent > Configuration**.

**Step 3** Select an application from the **Application** drop-down list.

 **NOTE**

If different applications have different collection settings, the collection settings applied to all applications will overwrite the collection settings applied to a specific application.



**Step 4** Click  to enable data collection.

 **NOTE**


This function is enabled by default. When you do not need to collect tracing and topology data of a specific application, disable this function to reduce resource usage.


**Step 5** Configure the sampling ratio according to [Table 14-1](#).

**Table 14-1** Sampling ratio

Sampling Ratio	Scenario	Advantage
Common sampling: One (the first) piece of data is sampled among N pieces of data.	Collection probes are deployed on services. If collection probes frequently collect tracing data, system performance may be affected. It is recommended that the resource usage of collection probes be minimized. You are advised to set common sampling based on service requirements.  Click  , enter an integer greater than or equal to 0, and then click  .  For example, for 35 pieces of data, if one piece of data is sampled among 20 pieces of data, the first and twenty-first pieces of data are sampled.	100% sampling can be set.


Sampling Ratio	Scenario	Advantage
Intelligent sampling: 100% sampling is performed if a transaction error occurs or the call latency is greater than the Application Performance Index (Apdex) threshold. Otherwise, one piece of data is sampled among 100 pieces of data. For details about how to set the Apdex threshold for a transaction, see <a href="#">15 Configuration Center</a> .	For the timeliness and integrity of monitoring, all abnormal or slow call data can be recorded regardless of the sampling ratio, without affecting system performance. You need to determine whether to enable intelligent sampling based on service requirements.	<ul style="list-style-type: none"> <li>• The performance reliability of collection probes is improved.</li> <li>• No abnormal or slow call data is missed.</li> </ul>

**Step 6** Click  to enable memory monitoring.

To prevent probes from affecting service performance in peak hours, enable memory monitoring. When the instance memory usage is excessively high, probes enter the hibernation state. You can also click  to set the duration and memory usage.


 **NOTE**

- Memory usage = Memory used by the Java process/Maximum available memory
- Maximum available memory: Use the smaller value between the available memory quota of the container and the maximum heap memory of the JVM. The maximum heap memory of the JVM is the value of `-Xmx`, which is 25% of the maximum available memory of the JVM by default.
- The memory usage during collection suspension must be greater than or equal to that during collection restoration.

**Step 7** Click  to enable the function of adding trace IDs to logs.



A trace ID uniquely identifies a tracing. When this function is enabled, the system adds trace IDs to log files. You can accurately search for logs based on trace IDs, such as `fffffffe1c08cab`, `fffffffe1c08cad`, and `fffffffe1c08cae`, as shown in the following figure.

```
02:56:04.027 [http-nio-8080-exec-2] [txid=fffffffe1c08cab] INFO [PersistenceRestController.java:99] - trying to find all products
02:56:06.030 [http-nio-8080-exec-10] [txid=fffffffe1c08cad] INFO [PersistenceRestController.java:99] - trying to find all products
02:56:40.168 [http-nio-8080-exec-4] [txid=fffffffe1c08cae] INFO [PersistenceRestController.java:99] - trying to find all products
```



**Step 8** Click  to enable SQL analysis.

When this function is disabled, the SQL data is not affected, but you cannot use this function.

**Step 9** Set the HTTP response codes to be ignored.

To quickly and accurately locate abnormal tracing, and prevent probes from misreporting normal tracing data, such as custom response codes, you can set the HTTP response codes to be ignored. Such codes will not be recorded in the error record table. Click , enter the HTTP response codes to be ignored, and click . If multiple HTTP response codes exist, separate them by commas (,).

**Step 10** Set the errors and exceptions to be ignored.

To quickly and accurately locate abnormal tracing, and prevent probes from misreporting normal tracing data, you can set the errors and exceptions to be ignored. Such errors and exceptions will not be recorded in the error record table. Click , enter the errors and exceptions to be ignored, and click . If multiple Java exception classes exist, separate them by commas (,). The default value is null.

----End



# 15 Configuration Center

---



## Setting Apdex Thresholds

**Step 1** Log in to the Application Performance Management (APM) console.

**Step 2** In the navigation pane, choose **Configuration Center**.



**Step 3** Select an application from the **Application** drop-down list.

**Step 4** Set Application Performance Index (Apdex) thresholds. For details about Apdex and Apdex threshold, see [Apdex](#).



- Click  next to **Topology Apdex Threshold (ms)**, enter the topology Apdex threshold, and click  to save the threshold.

 **NOTE**

The default topology Apdex threshold is 100 ms.

- Click  next to **Transaction Apdex Threshold (ms)**, enter the transaction Apdex threshold, and click  to save the threshold.

 **NOTE**

- The default transaction Apdex threshold is 500 ms.
- The setting takes effect for all application transactions. If an Apdex threshold has been separately set for a transaction, the currently set Apdex threshold takes effect for all transactions except this transaction. To separately set an Apdex threshold for a transaction, do as follows:
  1. In the navigation pane, choose **Transactions**.
  2. In the drop-down list in the upper left corner, select an application to which a transaction belongs.
  3. In the transaction list, click  under the **Apdex Threshold (ms)** column of the transaction, enter the current Apdex threshold, and click  to save the threshold.

----End

# 16 FAQs

[16.1 What Data Will Be Collected and What Are They Used For?](#)

[16.2 How Do I Obtain an AK/SK?](#)

[16.3 How Do I Obtain an AK/SK by Creating an Agency?](#)

## 16.1 What Data Will Be Collected and What Are They Used For?

When you enable data collection on Application Performance Management (APM), APM only collects service tracing data, resource information, resource attributes, memory detection information, and call request KPIs, and does not collect any privacy data. The collected data is used only for performance analysis and fault diagnosis, and is not used for commercial purposes.

Data Type	Collected Data	Transmission Mode	Storage Mode	Data Purpose
Tracing data	Tracing span data	HTTPS encryption and Access Key ID/ Secret Access Key (AK/SK) authentication for transmission	Project-based isolated storage	Query and display at the tracing frontend

<b>Data Type</b>	<b>Collected Data</b>	<b>Transmission Mode</b>	<b>Storage Mode</b>	<b>Data Purpose</b>
Call request KPIs	Call initiator address, receiver address, API, duration, and status	HTTPS encryption and AK/SK authentication for transmission	Project-based isolated storage	Calculation of transaction call KPI metrics (such as throughput, TP99 latency, average latency, and number of call errors), drawing of application topologies, and display of call metrics and topologies at the frontend.
Resource information	Service type, service name, creation time, deletion time, node address, and service release API	HTTPS encryption and AK/SK authentication for transmission	Project-based isolated storage	Query and display at the resource library frontend
Resource attributes	System type, system startup event, number of CPUs, service executor, service process ID, service pod ID, CPU label, system version, web framework, JVM version, time zone, system name, collector version, and LastMail URL	HTTPS encryption and AK/SK authentication for transmission	Project-based isolated storage	Query and display at the resource library frontend
Memory detection information	Memory usage, used memory, maximum memory, remaining memory, memory threshold-crossing time, and memory detection configurations	HTTPS encryption and AK/SK authentication for transmission	Project-based isolated storage	Query and display at the resource library frontend

## 16.2 How Do I Obtain an AK/SK?

### NOTE

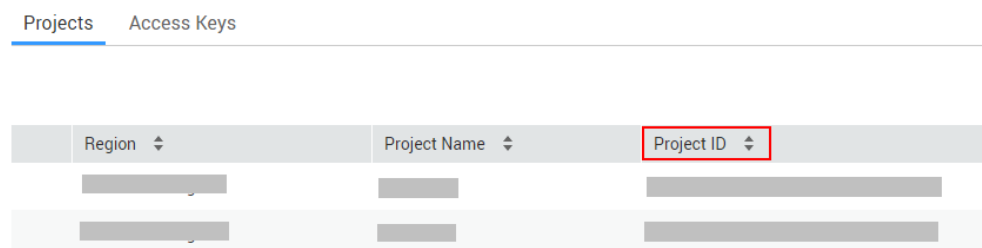
Each user can create a maximum of two Access Key ID/Secret Access Key (AK/SK) pairs. Once they are generated, they are permanently valid.

- AK: unique ID associated with the SK. It is used together with the SK to sign requests.
- SK: secret access key used in conjunction with an AK to sign requests cryptographically. It identifies a request sender and prevents the request from being modified.

### Procedure

- Step 1** Log in to the management console.
- Step 2** Hover over the username in the upper right corner and choose **My Credentials** from the drop-down list.
- Step 3** Obtain the project ID and AK/SK.
  - Obtain the project ID.  
On the **Projects** tab page, view the project ID in the **Project ID** column.

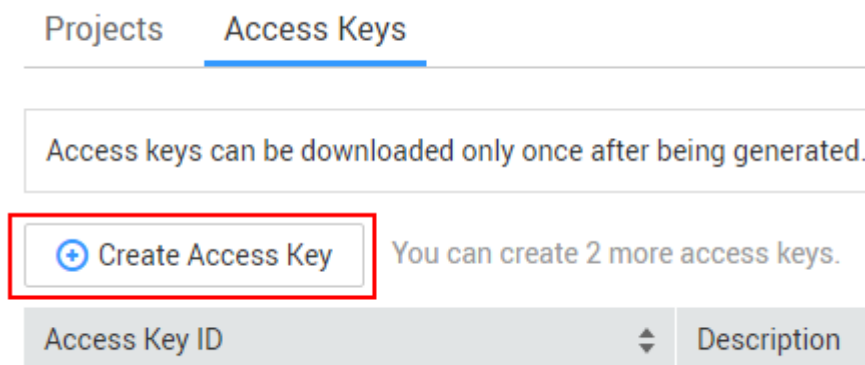
**Figure 16-1** Projects



Region	Project Name	Project ID

- Obtain the AK/SK.
  - On the **Access Keys** tab page, click **Create Access Key** to create an access key.

**Figure 16-2** Managing access keys



Projects **Access Keys**

Access keys can be downloaded only once after being generated.

**Create Access Key** You can create 2 more access keys.

Access Key ID	Description
---------------	-------------

- b. Enter the login password.
- c. Click **OK** to download an access key.

 **NOTE**

Keep the access key secure.

----End

## 16.3 How Do I Obtain an AK/SK by Creating an Agency?

After you create an agency, the ICAgent automatically obtains the Access Key ID/ Secret Access Key (AK/SK), helping you manage application performance.

### Creating an Agency

- Step 1** Log in to the Identity and Access Management (IAM) console.
- Step 2** In the navigation pane, choose **Agencies**.
- Step 3** On the page that is displayed, click **Create Agency** in the upper right corner. The **Create Agency** page is displayed.
- Step 4** Set parameters based on [Table 16-1](#).

**Table 16-1** Creating an agency

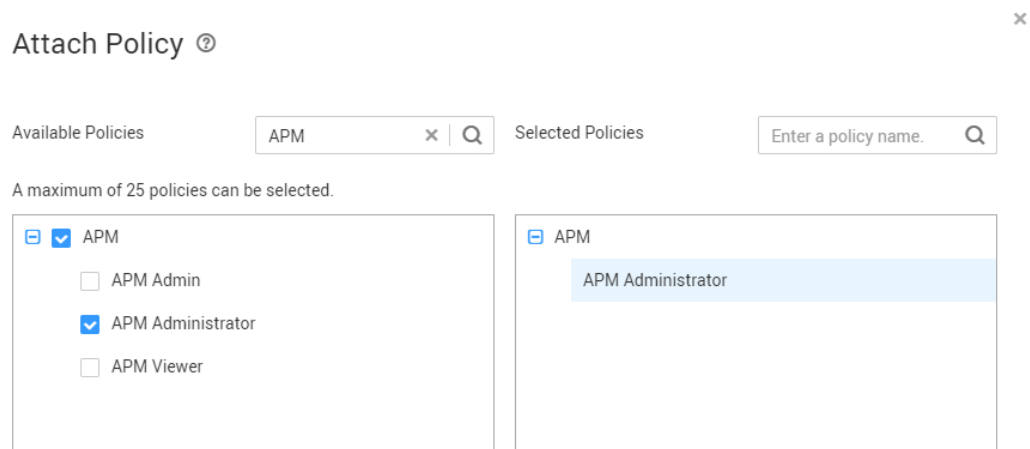
Parameter	Description	Example
Agency Name	Set an agency name.	aom_ecm_trust
Agency Type	Select <b>Cloud service</b> .	-
Cloud Service	Select <b>ECS BMS</b> .	-
Validity Period	Select <b>Unlimited</b> .	-
Description	(Optional) Provide detailed information about the agency.	-

----End

### Setting Permissions

- Step 1** In the **Permissions** area, click **Attach Policy** in the row of the current region. The **Attach Policy** page is displayed.
- Step 2** For **Available Policies**, enter **APM** in the search box on the left and select **APM Administrator** in the search result. In this way, the Application Performance

Management (APM) policy is synchronized to the **Selected Policies** area on the right.



**Step 3** Click **OK**. The agency relationship is successfully created.

----End

## Making an Agency Effective

**Step 1** On the management console, choose **Service List > Computing > Elastic Cloud Server**.

**Step 2** Click the ECS server where the ICAgent is installed. The ECS details page is displayed.

**Step 3** Select the created agency and confirm the configuration to make the agency effective.

**Step 4** (Optional) To set an agency for a newly created ECS, do as follows: On the **Create ECS** page, select **Configure now** for **Advanced Options** and select the created agency from the **Agency** drop-down list. Then, set remaining parameters and click **Apply Now**.

----End

# 17 Change History

---

**Table 17-1** Change history

Released On	Description
2020-02-26	This issue is the first release.